
Measuring and correcting wobble in large-scale transmission radiography

Thomas W. Rogers^{1,3} James Ollier² Edward J. Morton² Lewis D. Griffin^{1*}

¹Department of Computer Science, University College London, London, UK

²Rapiscan Systems Ltd., Stoke-on-Trent, UK

³Department of Security and Crime Sciences, University College London, London, UK

*Correspondence: l.griffin@cs.ucl.ac.uk

Abstract

BACKGROUND: Large-scale transmission radiography scanners are used to image vehicles and cargo containers. Acquired images are inspected for threats by a human operator or a computer algorithm. To make accurate detections, it is important that image values are precise. However, due to the scale of such systems, they can be mechanically unstable, causing the imaging array to wobble during a scan. This leads to an effective loss of precision in the captured image.

OBJECTIVE: We consider the measurement of wobble and amelioration of the consequent loss of image precision.

METHODS: Following our previous work, we use Beam Position Detectors (BPDs) to measure the cross-sectional profile of the X-ray beam, allowing for estimation, and thus correction of wobble. We propose: (i) a model of image formation with a wobbling detector array; (ii) a method of wobble correction derived from this model; (iii) methods for calibrating sensor sensitivities and relative offsets; (iv) a Random Regression Forest based method for instantaneous estimation of detector wobble, and (v) using these estimates to apply corrections to captured images of difficult scenes.

RESULTS: We show that these methods are able to correct for 87% of image error due wobble, and when applied to difficult images, a significant visible improvement in the intensity-windowed image quality is observed.

CONCLUSIONS: The method improves the precision of wobble affected images, which should help improve detection of threats and the identification of different materials in the image.

1 Introduction

Large-scale transmission radiography has become an essential tool for detecting threats inside vehicles and cargo containers. Threats may be related to customs (drugs, counterfeit goods, banned imports, stowaways, stolen cars) or security (firearms, improvised explosive devices, special nuclear materials, missiles) [26, 23, 25, 24, 8, 22]. Transmission radiography systems have become a mainstay of customs and border agencies around the world, and are finding increasing use in areas such as defence and the security of critical infrastructure, ports and events.

Images acquired by large-scale transmission radiography (Fig. 1) are inspected by a human operator or increasingly by computer algorithm [8, 18]. Detection of threats by operators is assisted by intensity manipulation (windowing, logarithms, histogram equalisation) and pseudo-colouring [3]. Additionally, scanners that acquire images at multiple photon energies permit material separation [14]

to be visualised based on differential absorption. On the basis of this visual inspection, the operator will either flag the vehicle for manual inspection or allow it to continue unimpeded.




Figure 1: A raw transmission X-ray image of a cargo container containing vehicles and vehicle parts (top) and an intensity manipulated version (bottom). Intensity manipulation is often used to reveal details in the image when searching for threats.

In order to detect threats, high spatial resolution and accurate image values are required [16]. The former, because threats may be small, and the latter because threats may be shielded by other cargo or only revealed by subtle differential absorption. State-of-the-art transmission systems offer imaging of vehicle contents at resolutions of a few mm/pixel [11] and precisions of 16 bits. In some systems, mechanical instability (wobble) leads to effective loss of precision. Whilst large-scale X-ray Computed Tomography (CT) could alleviate the issues of wobble and shielding, such systems are not widely deployed because they are too expensive and inefficient to be competitive [2, 27].

In our previous work [19] we proposed that wobble can be measured using Beam Position Detectors (BPDs) which are placed perpendicular to the imaging array (Fig. 2). Wobble was estimated by performing Gaussian model fitting to the BPD data to obtain instantaneous beam position estimates. These instantaneous estimates were Bayesian fused with an estimate from an Auto-Regression (AR) to make estimates more robust for scanning moments where the BPD was non-uniformly obscured by an object in the scanned scene. The wobble estimates were then used to make corrections to air-only images in order to quantify performance. We determined that we could correct 70% of image error due to wobble.

In this work, we follow a similar approach using BPDs but with several new contributions: (i) a model of image formation in the presence of wobble and other scanner design imperfections such as variable imaging sensor misalignments, variable sensor responses, and source fluctuation; (ii) improved wobble estimation using a Random Regression Forest (RRF) model for improved instantaneous estimation of wobble and its uncertainty, (iii) improved image correction by estimating the relative offsets of sensors, and (iv) estimation of sensor sensitivities by Sum of Squared Error (SSE) minimisation model fitting. Furthermore, we extend testing of image correction methods to include qualitative evaluation on images of complicated scenes.

In the next section we set out the technical background and review related work. In Sec. 3, we give a precise description of the effects of wobble and a method of image correction based on using BPDs to estimate fixed (e.g. sensor sensitivities, sensor offsets, and beam geometry) and dynamic (e.g. wobble and photon flux) system parameters. In Sec. 4 we propose a method for estimating the dynamic system parameters. Finally, in Sec. 5, we test these methods on images that we have collected from a Rapiscan Eagle®G60, a large-scale transmission X-ray gantry system, modified by the addition of four BPDs.




Figure 2: A typical transmission radiography system. Translation of the scene relative to the source and detector produces image columns, whilst each image row corresponds to a single sensor position in the imaging array. The set-up, considered in this paper, has been modified by addition of four Beam Position Detectors (BPDs) which are detector strips oriented at 90° to the imaging array. These allow the intensity profile across the beam width to be measured.

2 Background and Related work

Large-scale transmission radiography scanners operate either in portal or traverse mode, and are sometimes capable of both [17]. In portal mode the scanner is stationary and the scene moves between the source and imaging array at a controlled speed. In traverse mode the detector and source move either side of the stationary scene. Portal mode is most useful in high-throughput scenarios; vehicles can drive through the scanner arch without the driver having to exit the vehicle or a rail-scanner can scan multiple cargo containers carried by train at up to 60 km/h [18]. Traverse mode is useful in security scenarios where an unoccupied vehicle cannot be interfered with, such as if it suspected to be a car- or truck-bomb, or if it needs to be covertly inspected so as not to raise suspicions. The traverse mode is also useful for scanning lines of stationary cargo containers at ports [17]. The traverse mode has advantages in some cases: (i) the scanned vehicle is unoccupied, so higher doses can be used, resulting in higher precision images; (ii) there is greater control over scanning speed and detector-object distance resulting in less spatial warping of the captured image; and (iv) they have a compact scanning footprint [15].

In the traverse mode, the imaging array may wobble as it moves across the scene due to uneven ground or vibrations from the engine (truck systems), oscillations in the boom (truck and rail systems), or due to wind or vibrations from traffic (truck and rail systems). This has a particular impact when operators search for threats placed in dense scenes, since under intensity windowing [3] the wobble artefact becomes apparent (Fig. 3). Furthermore, wobble reduces the quality of material separation images [14] since their computation is dependent on precise values. Discrimination of high atomic numbers is particularly important as it can reveal smuggled nuclear materials, or their shielding [15, 4]. Wobble occurs in both truck-mounted and gantry systems. In truck-mounted systems wobble is variable from scan to scan, but in gantry systems it is systematic. In this work we study a gantry system, since it allows determination of the wobble ground truth, but our methods can equally be applied to truck-mounted systems.

To our knowledge, other than our previous work [19], there have been no publications on addressing wobble in large-scale transmission radiography. However, wobble leads to artefacts in a range of imaging devices, including micro-CT and C-arm CT. We describe the most relevant work here and how it relates to the wobble effect that we attempt to measure and correct in this work.

C-arm CT systems suffer from wobble as the gantry rotates. This means that individual projections are translated relative to those captured by an wobble-free ideal device. Authors note that the wobble of the C-arm gantry is often repeatable over periods of up to two years and so wobble artefacts can be corrected by a one-off system calibration [21, 6]. This is similar to some large-scale transmission




Figure 3: An X-ray image of a fork-lift truck from a mobile scanner with mechanical instability (left) and the same image with material discrimination applied (right). The image grey-levels have been windowed to make visible the small changes in image value due to wobble. Wobble leads to a rippling curtain effect across the image. Each rectangular test piece corresponds to a single material of uniform thickness. The wobble artefact effects the classification of material type; the classification of a single test piece can change from plastic through to steel due to wobble. This results in a colour change across the test piece (indicated by red arrow) in the material discrimination image, where there should be no change.

systems, particularly those that are in fixed deployment and the gantry moves along rails, where the wobble effect tends to be systematic. However, in truck-mounted mobile systems wobble is much more unpredictable due to variable scan speed and due to variations in the topology of the surface that the truck traverses. Moreover, in C-arm CT wobble artefacts tend to lead to a blurring effect in the reconstructed image due to the misalignment of individual projections, whereas in large-scale transmission systems, wobble mostly leads to image intensity variations as the fan-beam comes in and out of alignment with the detectors. Indeed, geometric image distortions can be observed if wobble is particularly severe, but this will be the focus of later work.

Silver *et al.* [21] propose a method for determining and correcting wobble in C-arm CT. The authors assume that the wobbling motion of the C-arm is the identical for each image capture process, and so calibrate wobble correction based on a phantom image. The phantom consists of a helical structure of tungsten carbide spheres (pellets). The calibration computes wobble coefficients that are used directly in image reconstruction to obtain a wobble-artefact free image. The wobble coefficients are determined by fitting a mapping from physical space to projection space using least-squares. Fahrig and Holdsworth [6] also adopt a calibration approach to determine projection translations. They use a bi-cubic spline interpolation to determine translated projections. Since the calibration process determines translations at discrete gantry angles, they linearly interpolate between them to obtain estimates for different projection angles if required.

Wobble is also observed in micro-CT systems, but the wobble manifests in the rotation table since the detector and source are kept stationary [20]. In this case, wobble again leads to a blurring effect in the image, quite different to the effect observed in large-scale transmission systems. Authors have investigated image-based, calibration and online methods to correct for wobble.

Sasov *et al.* [20] investigate and evaluate an image-based and a calibration-based method. The image-based method is an iterative compensation scheme, which first does an initial reconstruction using filtered back-projection, yielding blurry wobble affected reconstruction. Estimates for projection translations to compensate for wobble are determined by comparing original projections with corresponding forward-projected image estimates. The comparison is done either by cross-correlation or least-squares. Under these translations a new reconstruction is made and the process is iterated until the reconstructed image is satisfactory. The calibration-based method, measures wobble in a short reference scan directly before or after image capture to determine the compensatory translations of individual projections. They measure the position of the focal spot, relative to a metal pin placed in the scene, by fixing a fine metal mesh to the X-ray source. The authors claim that the second method is more suitable for slow drifts (wobble) and that the approach is faster and less computationally demanding than the iterative based method. However, the image-based method has the advantage of working purely on measured image values.

Zhao *et al.* [29] Propose an online method, which uses capacitive distance sensors to measure the wobble of the rotation table in Micro-CT. The measurements are used to translate individual projections to compensate for the displacement of the rotation table due to wobble. The authors report that the methods improve images by 53.1% and 65.5% when calibrating projections in the horizontal and vertical directions, respectively.

Due, to the unpredictable component (e.g. wind, uneven topology, vibrations) of wobble in large-scale transmission radiography, it is not possible to correct wobble purely by calibration. Image-based methods (without using BPDs or prior knowledge about large-scale radiography), such as Total Variation (TV) denoising or Translation Invariant Wavelet Shrinkage (TIWS) [12], may be applicable, however they are difficult to use in practice without prior information on the severity of the wobble artefact which we measure (online) in this work. In this contribution, we use both a calibration procedure and an online method. The calibration procedure is used to estimate a number of parameters that are fixed for a given system, including: misalignments of imaging sensors; the collimated width of the fan-beam; and the sensitivities of individual sensors due to housing attenuation and their intrinsic response. The online component, is for the estimation of wobble and estimation of the fluctuation in the photon flux, which can both vary unpredictably during a traverse mode scan. We describe these methods in the next section.

3 A model of image formation with wobble

To describe image formation with a wobbling detector, we use three coordinate systems (Fig. 4). We denote: the coordinates of imaging sensor pixels along the Γ -shaped imaging array (image vertical) by $y \in \mathbb{Y}$; the time coordinates indexing each scanning moment during image acquisition (image horizontal) by $t \in \mathbb{T}$; and the coordinates along the orientation of the BPDs (perpendicular to the beam and imaging array) by $x \in \mathbb{X}$. The origin $x = 0$ is taken as the vertical midline (dashed in Fig. 4) of the imaging array.

The formation process of an image $I_{ty} \in \mathbb{R}^+$ is described as follows. The X-ray source emits a photon flux $A_t \in \mathbb{N}$ at scanning moment t . This flux is collimated into a fan-beam of width β_y , which has a spatial distribution on the imaging plane according to

$$\exp(-(b_{ty} - d_y)^2 / (2\beta_y^2)). \quad (1)$$

The parameters $b_{ty} \in \mathbb{X}$ define the displacements of the beam cross section maximum from the vertical midline: when wobble occurs this varies with t and y , without wobble only with y . The parameters $d_y \in \mathbb{X}$ are the horizontal offsets of the imaging sensors from the vertical midline. For a given linear ID with endpoint offsets $\{\delta_l, \delta_u\}$ we constrain d_y to a linear function

$$d_y := (y_u - y_l)^{-1}((y_u - y)\delta_l + (y - y_l)\delta_u), \text{ where } y_l < y < y_u. \quad (2)$$

The X-ray photons pass through the scene and interact via absorption and scattering, and we denote the scene transmission by $S_{ty} \in [0, 1]$. This is dependent on the thickness and type of material composing the scene. The final measured image is determined according to a sensitivity factor $R_y \in [0, 1]$, which incorporates (i) the fraction of photons that are transmitted through the sensor housing and not absorbed or scattered, and (ii) the fraction of photons impinging on the detector that are counted (the intrinsic response of a sensor).

Therefore, the final image, assuming no cross-pixel effects such as photon scatter or detector cross-talk, is approximated by

$$I_{ty} = A_t \cdot \exp(-(b_{ty} - d_y)^2 / (2\beta_y^2)) \cdot S_{ty} \cdot R_y. \quad (3)$$

The scene transmission S_{ty} is the physical quantity that we are trying to measure, therefore the ideal image is

$$\underbrace{S_{ty}}_{\text{ideal}} = \underbrace{I_{ty}}_{\text{raw}} \cdot \underbrace{(A_t \cdot \exp(-(b_{ty} - d_y)^2 / (2\beta_y^2)) \cdot R_y)^{-1}}_{\text{correction factor}}. \quad (4)$$

To obtain the ideal image, one must estimate the different components of the correction factor. In the *portal* scanning mode, correction is straightforward. Absence of wobble means that $b_{ty} = b_y$, so that all that needs to be dealt with is:




Figure 4: *Left*: Part of the imaging array showing two misaligned Imaging Detectors (IDs), a Beam Position Detector (BPD), and a wobbling fan-beam. The magnitude of the wobble and the sensor misalignments have been exaggerated in this figure. The offsets d_y for individual imaging sensors are confined to a linear function determined by the offsets $\{\delta_l, \delta_u\}$ of the ID endpoints. The fuzzy bars illustrate the fan-beam incident on the imaging array. The Gaussian (width β_y and position b_{t_1y}) shows the profile of the fan-beam on the BPD. *Right*: A later time point $t_2 > t_1$. Due to wobble, the fan-beam has moved relative to the imaging array so that the intensity recorded by the imaging detectors has changed. This leads to an effective loss of image precision. Correction requires estimation of the beam displacements b_{ty} and the offsets d_y to be estimated. The b_{ty} , dense in t and y , can be interpolated from estimates dense in t but computed at the sparse y values where BPDs are located.

1. image column variations due to fluctuations of the photon source A_t ;
2. image row variations due to sensitivity R_y , and the fixed position and geometry of the beam $\exp(-(b_y - d_y)^2 / (2\beta_y^2))$;
3. image pixel variations due to Poisson variation in the number of photons that reach an imaging sensor.

The image column and row variations (1 & 2) can be corrected by normalising the columns and rows in the image respectively. In this work we do not attempt to correct for Poisson variation (3), however there are several denoising algorithms for Poisson-distributed noise [5, 10] in the literature. Note that Poisson variation can also be ameliorated by increasing the beam intensity or exposure time, but this has implications on safety and cost.

In the *traverse* scanning mode, where wobble *does* occur, the correction is complicated. The beam position b_{ty} now varies with t as well as y , and the imaging sensor offsets d_y must now also be estimated. These and the other parameters in the correction factor (Eq. 4) can be separated into two classes; (i) *system parameters* ($\beta_y, \{\delta_l, \delta_u\}$ and R_y) that are estimated in a one-off calibration which we describe below, and (ii) *dynamic parameters* (b_{ty}, A_t) that are estimated per time point (online). The source variation A_t is straightforward to address by taking an image patch from a single ID close to the source and averaging it over rows. In the remainder of this paper we work on A_t -corrected images. In Sec. 4 we describe a method to estimate b_{ty} .

In the one-off calibration, for each BPD we estimate β_y (beam width at the BPD location) and R_x (the sensitivity of the sensors along the BPD). For each ID, we estimate $\{\delta_l, \delta_u\}$ (the misalignments of the

ID at its endpoints) and R_y (the sensitivity of the sensors along the ID). The estimates are determined by model fitting to data collected during a traverse (wobbling) scan of an air-only scene. Although wobble has a detrimental effect on image precision, we benefit from wobble in these estimations since it allows us to disentangle (i) β_y and R_x , and (ii) b_{ty} and $\{\delta_l, \delta_u\}$.

The calibration is two-step and summarised as follows. First, we perform a Sum of Squared Errors (SSE) minimisation model fit using a Gaussian model of the fan-beam incident on the BPD, masked by the sensor sensitivities. In the fitted model, the Gaussian centre is allowed to vary freely with time but the beam width and sensitivities are unvarying. Having estimated the unvarying beam widths and the time-varying beam positions at each BPD, we linearly interpolate these to the positions of the sensors of the IDs. With these estimated, next we model fit to determine the ID parameters. We perform a SSE fit to the data from each ID to jointly estimate $\{\delta_l, \delta_u\}$ and R_y . The SSE is taken between the ideal image (raw image multiplied by correction factor, as in Eq. 4) and a uniform unit-valued image. The correction factor is composed using the interpolated β_y and b_{ty} estimates (from step 1), and the estimated parameters $\{\delta_l, \delta_u\}$ and R_y .

4 Wobble estimation algorithm

To estimate wobble for inhomogeneous scenes, we need to estimate b_{ty} at the BPDs, and then interpolate it along the IDs. However, the simple model fitting of the previous section is not applicable for inhomogeneous scenes. At some scanning moments the beam will be distorted from a Gaussian shape, and at other moments it will be undetectable due to dense loads. To cope with this, we estimate the beam position at time t by fusing an instantaneous estimate \hat{b}_{inst} (with uncertainty $\hat{\sigma}_{inst}$), with an estimate \hat{b}_{prior} (with uncertainty $\hat{\sigma}_{prior}$), based on the previous n beam position estimates.

4.1 Instantaneous estimation

The profile (D_{tx}) measured at each instant by a BPD, is a multiplicative combination of (i) the beam profile (P_{tx}), (ii) the scene transmission (S_{tx}), and (iii) the sensitivity (R_x). We estimate the beam profile from the measured profile, fixed estimates of the sensitivity (Sec. 3), and dynamic estimates of the scene transmission estimated from previous time-points of the BPD signal, according to

$$\hat{P}_{tx} = D_{tx} / (\hat{R}_x \hat{S}_{tx}). \quad (5)$$

This estimation works well in cases where the scene is not too dense (Fig. 5, 1b & 2b); but when it is the estimated beam profile can be inaccurate due to (i) the low (noise-dominated) sensor signal, or (ii) deviation of photon trajectories due to scatter (Fig. 5, 3b).

We estimate the scene transmission function \hat{S}_{tx} using measurements of the BPD as it slides across the scene. A given pixel on the BPD samples each point, at its y -value, in the scene (Fig. 6). Plotting the response of this pixel as a function of time gives an estimate of the scene transmission function. Since each of the BPD sensors also sample each point in the scene, we can construct a similar estimate for each sensor. The final estimate of \hat{S}_{tx} is obtained by taking a weighted average of the estimates from each of the sensors. We take the weighted average to reduce noise in the estimate from sensors that are aligned with the low signal tails of the Gaussian cross section.

With the estimate of the beam profile (\hat{P}_{tx}), we can estimate the instantaneous beam position b_{inst} and its uncertainty σ_{inst} . The estimator should be able to deal with non-linear relationships in the data and be able to produce data dependent uncertainty estimates. We have experimented with using Gaussian model fitting, as used in Ref. [19], but find that the non-normal distribution of the errors makes estimation of the uncertainty unreliable.

In this contribution, we use a Random Regression Forest (RRF) [1] to construct a robust estimator of the beam position from the beam profile estimates. A RRF model is based on an ensemble of decision trees and is capable of modelling non-linear relationships as required. Each tree in the RRF produces an estimate of the beam position. We obtain estimates of the instantaneous beam position \hat{b}_{inst} and its uncertainty $\hat{\sigma}_{inst}$ by taking the mean and standard deviation of the tree responses, respectively. We observe, for this study, that the standard deviation of the tree responses has a strong correlation with the actual error in the beam position estimate. Other advantages of RRF is that it is fast to train and deploy, and resistant to overfitting.

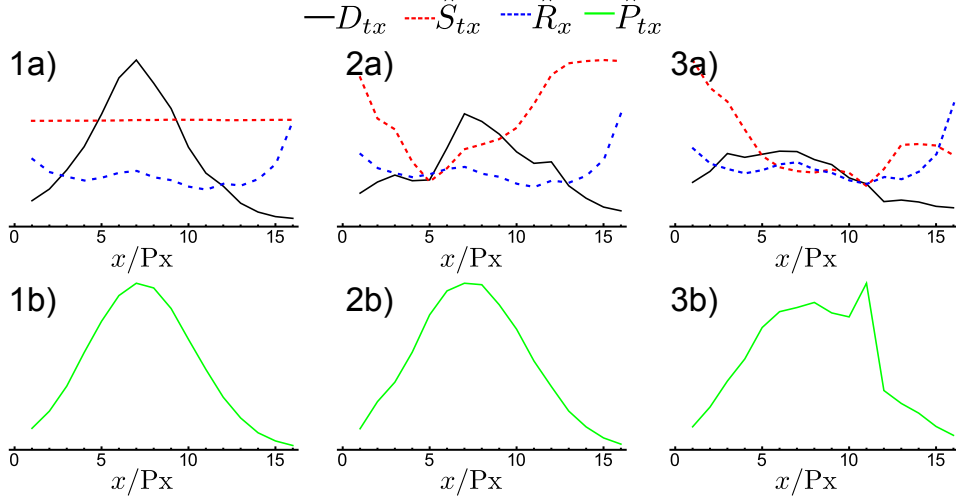


Figure 5: Examples of the estimated beam profile \hat{P}_{tx} (green) computed by dividing the measured Beam Position Detector (BPD) profile D_{tx} (black) by estimates of the scene transmission \hat{S}_{tx} (red) and the sensor sensitivity \hat{R}_x (blue). *Left:* Example of a homogeneous scene, and thus the estimate \hat{S}_{tx} is flat, resulting in a Gaussian \hat{P}_{tx} . *Middle:* Example of an inhomogeneous scene, the resulting \hat{P}_{tx} is approximately Gaussian. *Right:* Example of a dense inhomogeneous scene, where the resulting \hat{P}_{tx} is non-Gaussian which we attribute to photon scatter.




Figure 6: *Left:* Illustration of a Beam Position Detector (BPD) being translated across a scene during a scan. At consecutive timepoints $t = \{1, 2, \dots, T\}$ a given sensor (green) samples consecutive points in the scene. *Right:* plotting these samples as a function of t yields an estimate of the scene transmission. Each BPD sensor gives a similar estimate, and we perform a weighted average of them to reduce the noise in the final estimate \hat{S}_{tx} . Sensors towards the ends of the BPD, which receive low signal, are given a lower weighting in the average than those near the Gaussian centre which receive a higher signal.

In the RRF, N_t trees are constructed top-down with bagging and random subspace sampling. Internal nodes are split using standard thresholding, and optimised according to the Residual Sum of Squares (RSS). At each split m features (i.e. BPD pixels; elements of \hat{P}_{tx} at fixed t) are randomly sampled. For stopping criteria, we do not set a maximum tree depth and enforce a minimum of two samples per split. To tune the N_t and m , we first set m to the recommended default ($m = 1/3 \times \# \text{ features} = 5$) for regression. We then vary N_t and assess the RMSE to choose a sufficient number of trees so that the RMSE is stabilised but not too many that computation time is slow. With the N_t fixed, we then vary m from 3 to 12 to find the optimal RMSE, before verifying N_t again as before. By this method we determined that $N_t = 500$ was adequate and the default $m = 5$ was optimal.

In this work we use the `randomforest-matlab` implementation of RRFs [9]. For training, the ground truth values of the beam displacement were obtained by use of a gantry system in traverse

mode, described later in Sec. 5. We train a separate RRF for each BPD, using 1.4×10^5 measurements from five independent scans so that there is no overlap with the test images used in Sec. 5.

4.2 Estimation based on previous estimates

In cases where the BPD is heavily obscured (low signal-to-noise), the RRF-based instantaneous estimate will give a poor estimate of the beam position and a high uncertainty. In these cases, we want the beam position estimate to be sensible, and to achieve this we incorporate information about prior beam positions using an Auto-Regression (AR). The wobble of the detector array is partly deterministic (consider a swinging pendulum), but also stochastic due to the variable scanning surface, wind and vibrations. An AR is capable of learning some of the deterministic wobble whilst allowing for stochastic variation. It is also simple to implement and fast to compute. Additionally, we observe (Fig. 7) that the beam position trace has a high frequency component due to fluctuations of the photon source, possibly originating from electronic circuitry; and a low frequency component due to the wobble of the imaging array. The high frequency element makes simple estimation, based on the previous time point, unreliable. An AR, however, allows incorporation of n previous timepoints, where n can be tuned on data to achieve best performance. Moreover, the AR approach effectively smooths out erroneous estimates from previous time-points, but is beneficial over other smoothing filters (e.g. median filter) since it is possible to propagate previous errors to be used in fusion (Sec. 4.3) with the instantaneous estimate.




Figure 7: *Left:* Beam position b_{ty} as function of time t during a wobbling air scan. *Right:* A zoom so that the high frequency component of the wobble is visible.

The AR model predicts the current beam position based on a linear combination of the previous n beam positions with an added, normally distributed, perturbation

$$b_t = \sum_{t'} w_{t'} b_{t-t'} + N(0, \varepsilon^2) \text{ s.t. } \sum_{t'} w_{t'} = 1, \quad (6)$$

where $1 \leq t' \leq n$.

The Auto-Regression (AR) weights $w_{t'}$ are determined by model fitting Eq. 6 to an independent air scan. The constraint $\sum w_{t'} = 1$ ensures that the model does not have an unrealistic systematic drift. The uncertainty ε is determined by applying the model to a second air-only scan and computing the Root-Mean-Square Error (RMSE). The fitted model is used to generate the prior beam position estimate and its uncertainty according to:

$$\hat{b}_{\text{prior}} = \sum_{t'} w_{t'} \hat{b}_{t-t'}, \quad \hat{\sigma}_{\text{prior}}^2 = \sum_{t'} w_{t'} \hat{\sigma}_{t-t'}^2 + \varepsilon^2. \quad (7)$$

Note that the uncertainties from previous timepoints are propagated when forming this estimate, so that if the AR operates on previous estimates that are highly uncertain they are incorporated into the AR uncertainty, which is useful in the fusion step.

4.3 Fusion of estimates

To incorporate the information from the previous time-points, we fuse the estimates from the AR and RRF models according to their uncertainties. The fusion should weight the final estimate more towards the AR if the RRF-based estimate is more uncertain (e.g. due to low signal-to-noise).

Equally, if the AR uncertainty is high, because many of the previous n RRF-based estimates were also uncertain, but the next instantaneous estimate is very certain, then the fusion should weight more towards the RRF-based instantaneous estimate. To achieve, this we use a Bayesian fusion, which is equivalent to a Kalman Filter [7]. This approach is illustrated in Fig. 8.




Figure 8: Demonstration of Bayesian fusion of a prior estimate (black) and an instantaneous estimate (red) to obtain a fused estimate (blue). The width of the Gaussians correspond to the uncertainty on the estimate, and their centroid to the estimate value. The x -axis can be imagined as the Beam Position Detector (BPD). In (a) the prior has a higher certainty than the instantaneous estimate and so the fused estimate is weighted towards the prior, in (c) the opposite is true. In (b) both estimates have equal certainty and so the fused estimate compromises between the two.

To estimate the beam position \hat{b}_t and its uncertainty $\hat{\sigma}_t$, we Bayesian fuse the instantaneous estimate \hat{b}_{inst} and its uncertainty $\hat{\sigma}_{\text{inst}}$ (Sec. 4.1) with a prior estimate \hat{b}_{prior} and its uncertainty (Sec. 4.2). This is expressed as:

$$\hat{b}_t = (\hat{b}_{\text{inst}}\hat{\sigma}_{\text{prior}}^2 + \hat{b}_{\text{prior}}\hat{\sigma}_{\text{inst}}^2)/(\hat{\sigma}_{\text{prior}}^2 + \hat{\sigma}_{\text{inst}}^2), \text{ with } \hat{\sigma}_t^2 = (\hat{\sigma}_{\text{prior}}^2\hat{\sigma}_{\text{inst}}^2)/(\hat{\sigma}_{\text{prior}}^2 + \hat{\sigma}_{\text{inst}}^2). \quad (8)$$

This weights the two beam position estimates by their uncertainty. If the uncertainty of an estimate is low then that estimate contributes more to the fused estimate. In particular, if the instantaneous estimate is uncertain because of dense shielding, the prior estimate will be relied on; but when it is certain it will dominate the overall estimate.

5 Results

For the purposes of this study, and to test out our methods, we collected data using a modified Rapiscan Eagle®G60 transmission X-ray scanner. We rotated four of the IDs by 90° to become BPDs. The BPDs were placed at the extremes of the vertical boom and the horizontal boom, so that there were two BPDs per boom. The wobble characteristics are different at each location, for example wobble is most severe at the bottom of the vertical boom. Note that in a commercial implementation of BPDs, the system would have a full set of IDs with additional detectors for BPDs, but we have adopted this modification in experiments to reduce cost. We collected air-only images in portal and traverse modes, and several traverse mode scans of objects (e.g. trucks, forklifts, scissor lifts) were performed. The scanner operates at 90Hz and has a pixel size of 5.6mm, giving an effective spatial resolution of roughly 3mm. The scanner uses a Bremsstrahlung beam with a cut-off energy of 6MeV. This is the same energy used in commercial systems, and gives enough penetration to achieve reasonable signal-to-noise ratio on the BPD for most objects.

We adopted a gantry set-up, since it provides a ground truth for wobble. Wobble is observed in both gantry and truck-mounted systems, with a similar amplitude and frequency composition. However, for a gantry system, wobble is the same (modulo alignment) for each scan, but variable for truck-mounted systems. The gantry system allows us to obtain an accurate ground truth by aligning wobble estimates from an air-scan with the air parts of an object scan.

5.1 System parameter estimation

The system parameters β_y (beam width), R_y (sensitivities) and d_y (imaging sensor offsets) were estimated according to Sec. 3, and are shown in Fig. 9. Small and large y -values correspond to the bottom and top of the image, or the vertical and horizontal parts of the Γ -shaped imaging array, respectively. The gaps in y -values are where an ID has been removed or rotated to form a BPD.




Figure 9: Estimated system parameters: (a) beam width, β_y ; (b) sensor sensitivities, R_y ; and (c) horizontal imaging sensor offsets from the vertical, d_y . The dashed horizontal line marks the transition from the vertical (below) part of the Γ -shaped imaging array, to the horizontal (above). The black dots indicate the y -positions of the Beam Position Detectors (BPDs). Gaps in y -values are where an Imaging Detector (ID) has been removed or rotated to form a BPD in the experimental set-up. The beam width increases (decreases) as the distance from source to the array increases (decreases), due to dispersion. The sensitivities fluctuate between adjacent sensors due to their different intrinsic responses. The estimated sensor offsets are piece-wise linear because they are grouped by ID, and are of the order of a few mm which is within the manufacturing tolerance of a system of this scale.

The estimate of β_y (Fig. 9.a) increases as you go along the horizontal of the Γ -shaped imaging array and away from the source due to beam dispersion; it then decreases as you go along the vertical of the Γ -shaped imaging array and slightly closer to the source. The sensitivities R_y (Fig. 9.b) have a lot of variation between adjacent imaging sensors due to their intrinsic response and due to variations in the housing of the Γ -shaped array. The estimated offsets of the IDs (Fig. 9.c) are of the order of a few mm, which when compared to their 10cm length is plausible for a human engineer placing them during the construction of the scanner, and is indeed within the manufacturing tolerance of a scanning device of this scale (6m tall). Note that the piecewise-linear nature of d_y is due to the linear constraint places on each ID (Eq 2).

5.2 Wobble estimation

The AR was trained on an air-only traverse mode image. Fig. 10(a) shows the RMSE performance of the trained AR on an independent air-only test image as a function of the number (n) of previous timepoints considered. As n is increased the RMSE decreases, reaching a minimum at $n = 64$, before the RMSE begins to grow. When n gets too large the model overfits and performance deteriorates. We choose $n = 32$ since the RMSE is near optimal but requires half the number of parameters. The AR weights for $n = 32$ are shown in Fig. 10(b). It shows that more importance is placed on the most recent b estimates as expected. The oscillating structure is the AR system's way of coping with the high frequency component of the beam movement.

To assess the performance of the proposed beam position estimates, we test performance on “easy”, “intermediate” and “difficult” scenarios from the collected data. For each, we compare the new RRF-based method for instantaneous estimation to the old Gaussian-based method from Ref. [19]. We also compare the RRF-based method instantaneous method, with the fused estimate which we refer to as RRF-AR.

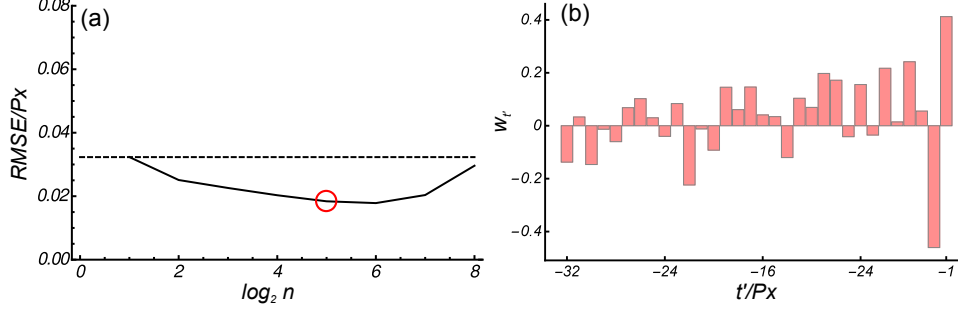


Figure 10: Auto-Regression (AR) model fit: (a) The Root-Mean-Square Error (RMSE) performance of the AR for different numbers of previous time-points n included in the model (dashed line indicates the standard deviation of the beam position, red circle indicates the near-optimal $n = 32$); and (b) the learnt AR weights $w_{t'}$ when $n = 32$. The RMSE decreases as the number of time points n included in the model increases, it reaches an optimum at around $n = 64$, before rising again due to overfitting. The AR weights have larger magnitude for the most recent time-points ($t' = -1, -2$) as expected, because these are most informative for predicting the next beam position. The oscillating structure in the weights is the AR’s way of coping with the high frequency wobble component.

For the “easy” scenario (Fig. 11), the RRF instantaneous estimate (green) is mostly accurate, with most estimates close to the groundtruth (black). The old Gaussian-based method (red) gives wildly inaccurate estimates when the BPD is occluded by an object thus resulting in a non-Gaussian BPD profile. However, the RRF yields estimates much closer to the groundtruth, in these cases. These estimates are made very accurate when fused with the AR (blue), since the RRF trees give variable responses which results in a larger uncertainty, so the fusion gives more weight to the AR. In particular, in Fig. 11(d) the fused estimate is much closer to the groundtruth than the RRF on its own.

In the “intermediate” scenario (Fig. 12). The old Gaussian-based method does even worse, and again the RRF-based method appears relatively robust to non-Gaussian BPD profiles, where the Gaussian-based method fails. In this scenario, fusion with AR, does not give a large change in estimates over just using the RRF since the RRF trees are confident in their estimate; there is not a large amount of variability in their votes. However, an improvement is seen in Fig. 12(d).

For the “difficult” scenario in Fig. 13, the old Gaussian-based method does even worse. The RRF-based instantaneous estimator appears far more robust, with estimates much closer to the groundtruth, however the performance is not as great as in the “easy” and “intermediate” scenarios. The fused estimates exhibit a bias (see 13.f) where the RRF performs poorly over a long time period. This happens where the total signal on the BPD is close to the background noise level (it is very heavily occluded by a truck engine), and hence the RRF finds it difficult to make accurate estimates of the beam position. This is reflected in the RRF uncertainty, and so the fused estimate puts full weight on the AR estimate, which results in a constant fused beam position estimate until a good instantaneous estimate is achieved. So the AR has forced the fused estimate into giving sensible estimates. Since the BPD signal is so low in this object and it occupies a large number of time-points, we reason that it would be impossible to obtain an accurate instantaneous estimate by any method based on the current BPD set-up.

For each of the scenarios (easy, intermediate, and difficult), we have quantified the performance of the methods in terms of: accuracy; bias; precision; and Mean Absolute Error (MAE) for the worst 5%, 1%, and 0.1% of time-points. We include the the worst MAEs since particularly bad time-points can be lost in the accuracy, precision, and bias metrics, particularly if there are many air-only time-points where estimation is straightforward. Moreover, wildly inaccurate wobble estimates could lead to column artefacts in the image after correction so are undesirable. The results are given in Table 1. For the “intermediate” and “difficult” scenarios, the RRF-based instantaneous estimation offers roughly an order-of-magnitude improvement across all metrics, over the old Gaussian-based method. For the “easy” scenario, this improvement is approximately 3-fold; the Gaussian method is already quite good at dealing with simple objects. By fusing the RRF with the AR (RRF-AR), the performance increases across most metrics, particularly for worst MAEs, however, there is little change (or a slight worsening for the “intermediate” scenario) in the overall accuracy. In the “easy” scenario there is

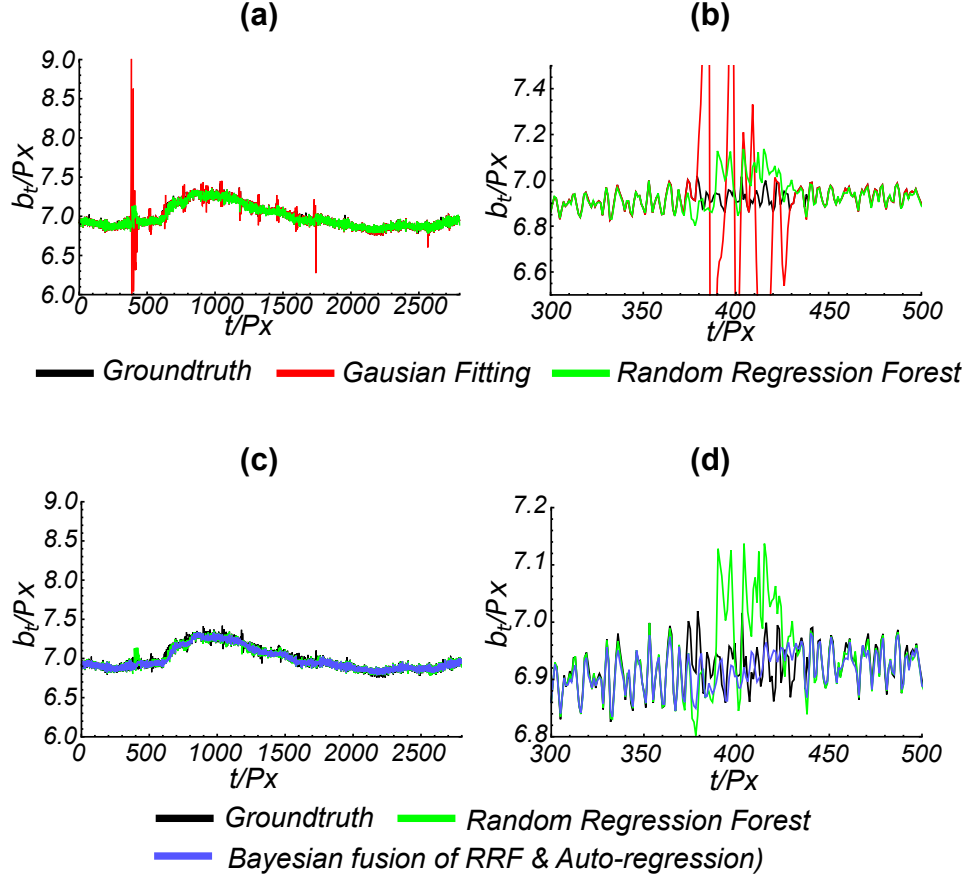


Figure 11: Beam position estimates for the “easy” scenario. In (a) the new Random Regression Forest (RRF) based method (green) for instantaneous estimation is compared to the old Gaussian-based method (red), and the groundtruth (black). In (c) the RRF-based method (green) for instantaneous estimation is compared to its Bayesian fusion with an Auto-Regression (RRF-AR; blue), and the groundtruth (black). Plots (c & d) show zooms for the most difficult region. The old Gaussian-based method gives wildly inaccurate estimates ($t/P_x = 400$) where the BPD profiles are occluded and so measure a non-Gaussian profile. The RRF yields much more accurate estimates, and is improved further (relative to the groundtruth) when fused with the AR (see d).

roughly a 15% improvement in the MAE for the worst 5% of time-points. For the “intermediate” case the improvement drops so about 5%. Finally, for the “difficult” scenario the worst 1% MAE improves by about 3%.

5.3 Image correction

We first assess the image correction method on an air-only scene. For air-only images, wobble estimation is straightforward, since the BPD profile is not distorted by obscuring objects in the scene. However, air-only images allow us to visualise and fully quantify the improvement from wobble correction. We can assess image quality based on the fact that a perfect (normalised) transmission air image would have all pixel values equalling unity. Image precision can therefore be assessed by computing the Root-Mean-Square (RMS) deviation or Peak Signal-to-Noise Ratio (PSNR) from this ideal.

In Fig. 14, we show air-only images from traverse and portal mode scans and their full correction split into stages. The stages are: sensitivity R_y correction (Fig. 14.b&f); wobble and ID offset $\exp(-(b_{ty} - d_y^2)/2\beta_y^2)$ correction (Fig. 14.c&g); and source variation A_t correction (Fig. 14.d&h). Images have been intensity windowed so that the wobble effect is visible in (Fig. 14.f). Note the visible difference between images (Fig. 14.b) and (Fig. 14.f), this difference is mostly due to wobble.

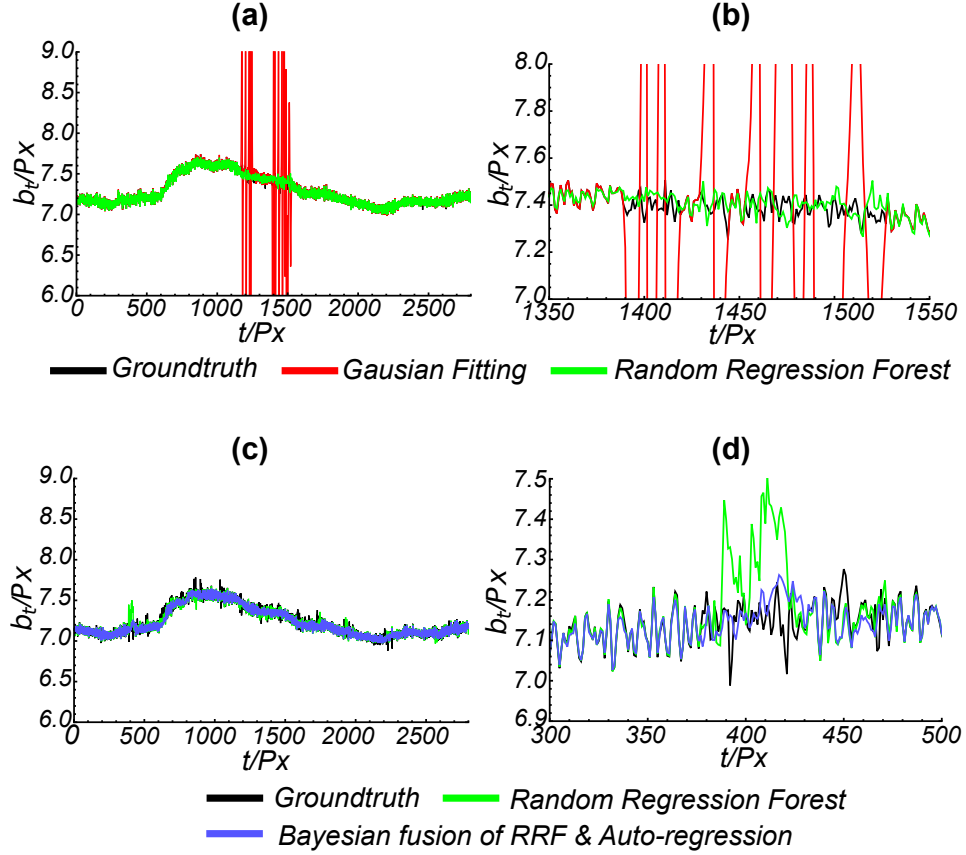


Figure 12: Beam position estimates for the “intermediate” scenario. In (a) the new Random Regression Forest (RRF) based method (green) for instantaneous estimation is compared to the old Gaussian-based method (red), and the groundtruth (black). In (c) the RRF-based method (green) for instantaneous estimation is compared to its Bayesian fusion with an Auto-Regression (AR; blue), and the groundtruth (black). Plots (b & d) show zooms for the most difficult regions. In this case, the old Gaussian-based method behaves very erratically ($t/Px = [1100, 1600]$), the RRF method gives estimates much closer to the groundtruth. However, some RRF-estimates are inaccurate (d), but these are improved when fused with the AR since the RRF trees give a larger uncertainty than the AR prior estimate.

Table 1: Performance metrics for: (i) the old Gaussian (Gauss.) based method of instantaneous estimation; (ii) the proposed Random Regression Forest based method for instantaneous estimation (RRF); and (iii) the Bayesian fusion of the RRF estimates with an Auto-Regression (RRF-AR). The metrics computed are: Accuracy (Acc.); Bias; Precision (Prec.); and the Mean Absolute Error (MAE) for the worst 5%, 1% and 0.1% of estimates. RRF gives an order-of-magnitude improvement over Gauss. for easy and intermediate scenarios, and 3-fold for the difficult scenario. RRF-AR gives 3-15% improvement in MAE depending on the difficulty.

Scenario	Meth.	Acc.	Bias	Prec.	5%	1%	0.1%
Easy	Gauss.	0.105	0.003	0.105	0.300	0.818	2.085
	RRF	0.033	0.004	0.033	0.122	0.168	0.203
	RRF-AR	0.030	0.006	0.029	0.104	0.140	0.167
Intermediate	Gauss.	0.470	-0.010	0.470	1.751	3.523	5.846
	RRF	0.019	-0.001	0.019	0.073	0.125	0.170
	RRF-AR	0.021	-0.001	0.021	0.069	0.111	0.149
Difficult	Gauss.	0.637	-0.135	0.623	2.012	3.329	5.670
	RRF	0.052	-0.008	-0.052	0.188	0.262	0.310
	RRF-AR	0.052	-0.014	0.050	0.191	0.253	0.284

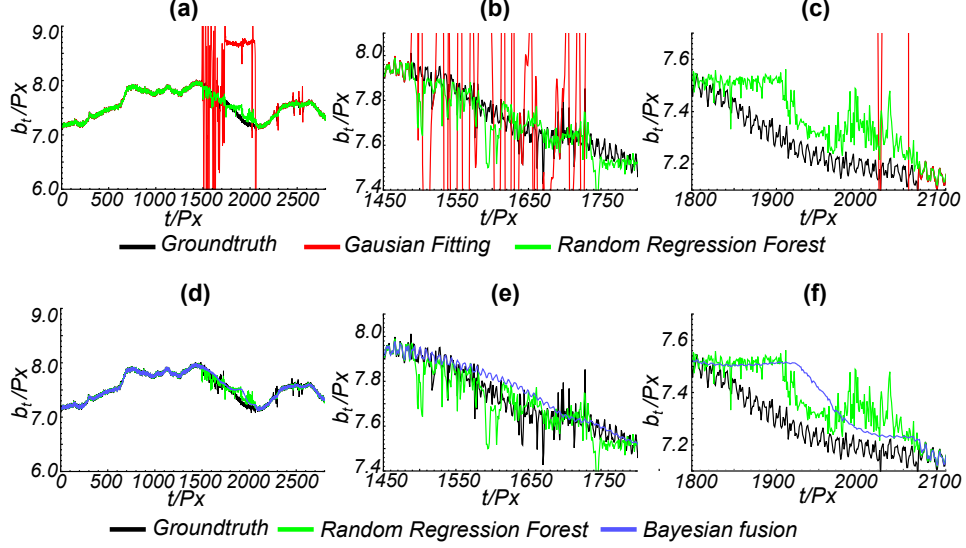


Figure 13: Beam position estimates for the “difficult” scenario. In (a) the new Random Regression Forest (RRF) based method (green) for instantaneous estimation is compared to the old Gaussian-based method (red), and the groundtruth (black). In (d) the RRF-based method (green) for instantaneous estimation is compared to its Bayesian fusion with an Auto-Regression (AR; blue), and the groundtruth (black). Plots (b, c, e, & f) show zooms for the most difficult regions. The RRF struggles to give accurate estimates in (c), because the BPD is passing across a truck engine, which is very dense and therefore the signal-to-noise is very low. This increases the RRF uncertainty, and so the fused estimate puts full weight on the AR estimate, which results in a constant estimate (f) until a better RRF estimates are achieved. So the AR has forced the Bayesian fusion into giving sensible estimates. This also occurs in (b) and (f) but to a much lesser extent.

The PSNR drops from 109dB to 77.2dB, from portal image (Fig. 14.b) to traverse image (Fig. 14.f) due to the wobble artefact. After wobble correction, to obtain image (Fig. 14.g), most of the wobble artefact is visibly improved. Indeed, the wobble correction improves the PSNR by 21.3dB but is unable to achieve the portal mode PSNR.

To make quantitative assessment of the effects visible in Fig. 15, the RMS deviations of the traverse and portal mode air-only images, before and after the different corrections, were used to deduce the magnitude of the noise sources before and after correction. Table 2 shows that wobble increases overall image noise, and has also reduced our ability to correct for *sensor sensitivity*, *ID offset*, and *source variation*. Although it is possible to correct for 99% of *sensor sensitivity*, the magnitude of *sensor sensitivity* is so large that it is still the second most dominant source of noise in the corrected image. Source variation was the least successfully corrected and this is apparent in Fig. 15, since the corrected images (Fig. 14.d&h) have some slightly visible column artefacts. Finally, we are able to correct 87% of wobble, thus outperforming our previous work [19], which did not incorporate sensor offset estimates into the correction.

The results for corrections applied to traverse mode images of a scissor lift and a forklift truck are shown in Fig. 15. Images have been intensity windowed, to the same range, to make the wobble artefact visible. The wobble correction is obtained using the Bayesian-fused beam position estimate. The red boxes indicate image regions most effected by wobble, and the green boxes show the same regions but after wobble correction. There is a visible improvement in the wobble artefact after wobble correction, showing that a good level of correction is obtained even when the BPDs pass through dense objects such as a fork-lift truck.

Fig. 16 shows image corrections on a truck image. Since the truck occupies most of the image, it is more difficult to see the effects of wobble and the corrections. The most obvious places are the steps up to the driver’s cabin and the area surrounding the test object. These are indicated by the red boxes in Fig. 16.b. After wobble correction (green boxes in Fig. 16.c), the artefact is reduced so that the driver’s steps and the test object become visible. In plots Fig. 16.i&ii we plot a column and row of

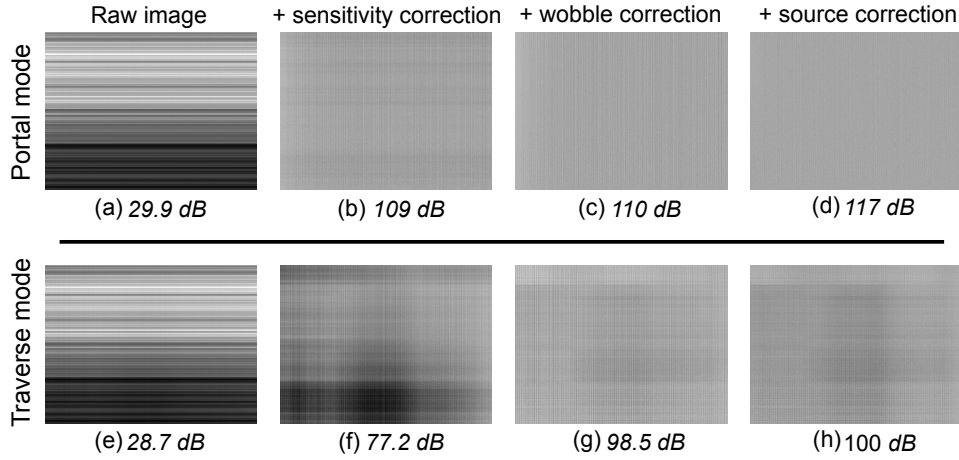


Figure 14: Images of portal (a-d) and traverse (e-h) mode air-only scans at different stages of correction, including: the raw image; the image after correction for sensor sensitivities; the image after correction for wobble (and sensor offsets); and the final image after after source correction. Corrected images have been intensity windowed so that the wobble artefact is visible in (b). For each image, the Peak Signal-to-Noise Ratio (PSNR) is given in decibels (dB). A noiseless and artefact-free air-only image should be uniform. The wobble artefact is clearly visible in the traverse image after the sensitivities have been corrected (f), and it is not visible in the portal mode image (b) since this mode is not effected by wobble. The PSNR is reduced by 31.8dB by the wobble artefact. After wobble correction (g) there is a visible improvement in the artefact, and the PSNR has improved by 21.3dB.

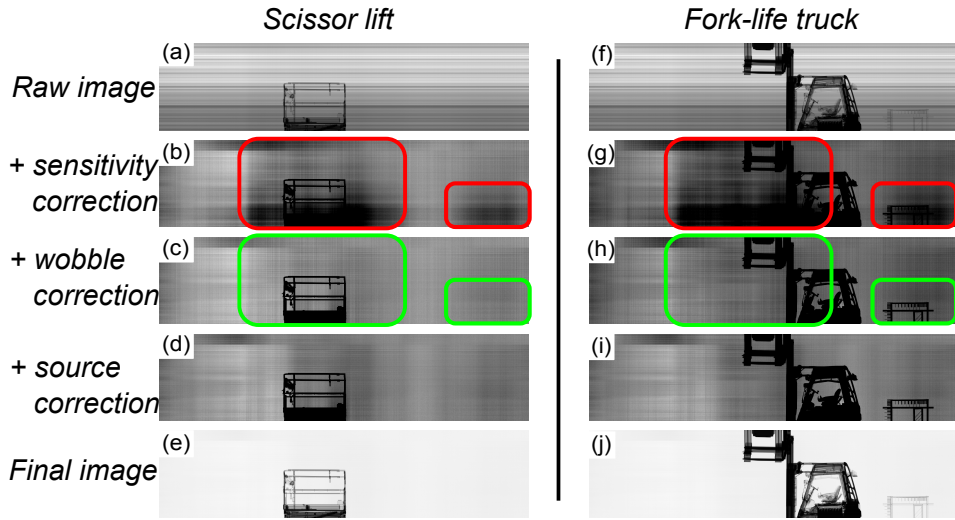


Figure 15: Images from traverse mode scans of a scissor lift (a-e) and a fork-left truck (f-j), after a series of corrections for: sensor sensitivities (b & g); wobble and Imaging Detector (ID) offsets (c & h); and source variation (d & i). Corrected images have been intensity windowed so that the wobble artefact is visible in (b) and (g). The final images (e & j) are the non-windowed versions of (d & i). The red boxes indicate regions where wobble is particularly visible, and the green boxes indicates the same regions after wobble correction. There is a clear visible improvement in the wobble artefact after correction, and so wobble measurement and correction works quite well even when BPDs are heavily occluded by dense object.

Table 2: RMS deviation contributions from different noise sources before and after corrections for: sensor sensitivity; Imaging Detector (ID) offsets; wobble; and source fluctuation. We do not attempt to correct Poisson noise in the photon counts.

Scan mode	Noise source	Symbol	Before	After	Reduction
portal	<i>sensor sensitivity</i>	R_y	0.2305	0.0000	100%
	<i>offset of ID endpoints</i>	$\{\delta_l, \delta_u\}$	0.0013	0.0000	100%
	<i>wobble</i>	b_{ty}	0.0000	0.0000	—
	<i>source variation</i>	A_t	0.0030	0.0000	100%
	<i>photon count</i>	—	0.0029	0.0029	0%
traverse	<i>sensor sensitivity</i>	R_y	0.2305	0.0026	99%
	<i>offset of ID endpoints</i>	$\{\delta_l, \delta_u\}$	0.0013	0.0004	72%
	<i>wobble</i>	b_{ty}	0.0185	0.0054	87%
	<i>source variation</i>	A_t	0.0030	0.0004	74%
	<i>photon count</i>	—	0.0029	0.0029	0%

pixels, respectively. In each, the red plot is from Fig. 16.b before wobble correction, and the blue plot is from Fig. 16.c after wobble correction. The pixels are taken from image lines that should have approximately constant (or piece-wise constant) pixel values. However, due to the wobble artefacts they are distorted from constancy. The wobble correction corrects a large part of this distortion.




Figure 16: Images from a traverse mode scan of a truck, after a series of corrections for: sensor sensitivities (b); wobble and Imaging Detector (ID) offsets (c); and source variation (d). Corrected images have been intensity transformed so that the wobble effect is visible in (b). The final image (e) is the non-windowed version of (d). The red boxes indicate regions around the driver's steps and a test object, where wobble is particularly visible. The green boxes indicates the same regions after wobble correction and one can see a visible improvement in the wobble artefact. The plots (i) and (ii) show plots of the pixel intensities across a column and row in the image, respectively. The red traces are uncorrected for wobble and taken from (b), whilst the blue traces are corrected for wobble and taken (c). The red traces should be approximately (piece-wise in ii) constant, however they are distorted by wobble. The wobble correction corrects most of this distortion.

6 Conclusion

We have proposed a series of image corrections to ameliorate detector wobble artefacts in large-scale transmission radiography. The corrections were derived by considering a model of X-ray image formation in the presence of a wobbling detector. The correction relies on the estimation of a number of fixed system parameters and dynamic parameters which vary during a scan. The fixed parameters include sensor sensitivities, sensor misalignments, and the width of the X-ray fan-beam. The dynamic

parameters include the position of the beam at different points along the detector array, and the fluctuation of the number of photons emitted by the source. We proposed a method for estimating the fixed system parameters by model fitting to an air calibration image.

Wobble is more difficult to estimate, and we adopt a similar approach, using Beam Position Detectors (BPDs), to our previous work [19]. BPDs are placed perpendicular to the imaging array, and measure the cross-sectional profile of the photon beam after interaction with the scene, allowing the position of the beam to be determined and hence detector wobble to be measured. In our previous work, we measured wobble by fitting a Gaussian model to the beam profile to extract an instantaneous estimate of the beam position. This was Bayesian-fused with a prior estimate based on an Auto-Regression (AR). In this contribution, we proposed a new instantaneous estimator based on a Random Regression Forest (RRF). We first estimate the true beam profile, as if the beam had not been attenuated by the scene, and then estimate the beam position and its uncertainty by taking the mean and standard deviation of the responses from a RRF, respectively.

To test the wobble estimation and image correction methods, we collected image data of several objects ranging in difficulty from a small scissor lift to a large truck. We used a commercial scanner, which we modified by rotating four imaging detectors by 90° to act as BPDs. Our new RRF-based approach to instantaneous estimation performs significantly (an order of magnitude in most cases) better than Gaussian fitting [19]. Moreover, its fusion with an AR achieves results close to ground truth, even for difficult objects, and performs better than the RRF by 3-15% in the worst cases. It struggles for cases where the object has a low signal-to-noise ratio for long durations in the scan, and we believe that this problem cannot be solved by wobble estimation based solely on BPD readings, unless one can accurately predict future beam positions from a limited number of accurate prior position estimates. This is unlikely due to the stochastic nature of wobble originating from uneven scanning surfaces or wind. Incorporation of measurement devices, such as accelerometers placed along the imaging array, may improve estimates even where there is almost no BPD signal due to object occlusion. This will be a focus of future work.

The wobble and system parameter estimates were used to apply corrections to images. We applied corrections to traverse and portal mode air-only images and achieved a reduction of 87% of image error due to detector wobble, thus improving on our previous work [19]. The wobble correction method was also applied to difficult images of objects and a notable qualitative improvement in the intensity-windowed image quality was observed, clarifying dense regions of the scene and mitigating human error. The method should also allow for improved material discrimination in images captured from dual-energy scanners in traverse mode. State-of-the-art material discrimination, for cargo, is performed by taking the log-ratio (or difference) of images at different energies, and relies on subtle differences between the images [14, 28, 13]. But in commercial traverse-mode systems material discrimination is often inaccurate due to image noise, including from wobble (Fig. 3). And so wobble correction as pre-processing step could help improve material discrimination accuracy. Testing this, and fully quantifying the effect of wobble on material discrimination, will be left to future work.

Future work, will include experimenting with other measurement devices, such as accelerometers, to improve the prior estimate of the beam position in cases where the RRF fails to obtain accurate estimates over long time-periods due to large, dense objects such as a truck engine. Potentially, beam position estimation could be improved by using more BPDs or even a 2D imaging array, and this will be explored. Additionally, we will investigate the severity and correction of geometric distortions cause by extremely heavy wobble. Such distortions can cause straight lines to become wobble, which potentially impacts on the performance of human operators searching for threats, particularly if their shape is distorted in an unnatural way.

Acknowledgements

Funding for this work was provided through the EPSRC Grant no. EP/G037264/1 as part of UCL's Security Science Doctoral Training Centre, and Rapiscan Systems Ltd.

7 References

- [1] L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.

- [2] N. Calvert, E. J. Morton, and R. D. Speller. Preliminary monte carlo simulations of linear accelerators in time-of-flight compton scatter imaging for cargo security. *Crime Science*, 2(1):1, 2013.
- [3] G. Chen. Understanding X-ray cargo imaging. *Nuclear Instruments and Methods in Physics Research Section B*, 241(1):810–815, 2005.
- [4] G. Chen, G. Bennett, and D. Perticone. Dual-energy X-ray radiography for automatic high-Z material detection. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 261(1):356–359, 2007.
- [5] W. Cheng and K. Hirakawa. Minimum risk wavelet shrinkage operator for poisson image denoising. *IEEE Transactions on Image Processing*, 24(5):1660–1671, 2015.
- [6] R. Fahrig and D. Holdsworth. Three-dimensional computed tomographic reconstruction using a C-arm mounted XRII: image-based correction of gantry motion nonidealities. *Medical Physics*, 27(1):30–38, 2000.
- [7] R. Faragher. Understanding the basis of the Kalman filter via a simple and intuitive derivation. *IEEE Signal Processing Magazine*, 29(5):128–132, 2012.
- [8] N. Jaccard, T. W. Rogers, and L. D. Griffin. Automated detection of cars in transmission X-ray images of freight containers. In: *Proceedings IEEE International Conference on Advanced Video and Signal Based Surveillance*, 387–392, 2014.
- [9] A. Jaialtil. Classification and regression by randomforest-matlab. Available at: <http://code.google.com/p/randomforest-matlab>, 2009.
- [10] S. Lefkimmiatis, P. Maragos, and G. Papandreou. Bayesian inference on multiscale models for poisson intensity estimation: Applications to photon-limited image denoising. *IEEE Transactions on Image Processing*, 18(8):1724–1741, 2009.
- [11] Y. Liu, B. D. Sowerby, and J. R. Tickner. Comparison of neutron and high-energy X-ray dual-beam radiography for air cargo inspection. *Applied Radiation and Isotopes*, 66(4):463–473, 2008.
- [12] A. Mouton, G. T. Flitton, S. Bizot, N. Megherbi, and T. P. Breckon. An evaluation of image denoising techniques applied to ct baggage screening imagery. in: *Proceedings IEEE International Conference on Industrial Technology*, 1063–1068, 2013.
- [13] S. Ogorodnikov, M. Arlychev, I. Shevelev, R. Apevalov, A. Rodionov, I. Polevchenko, et al. Material discrimination technology for cargo inspection with pulse-to-pulse linear electron accelerator. In: *Proceedings International Particle Accelerator Conference*, 2013.
- [14] S. Ogorodnikov and V. Petrunin. Processing of interlaced images in 4–10 MeV dual energy customs system for material recognition. *Physical Review Special Topics – Accelerators and Beams*, 5(10):104701, 2002.
- [15] V. J. Orphan, E. Muenchau, J. Gormley, and R. Richardson. Advanced γ ray technology for scanning cargo containers. *Applied radiation and Isotopes*, 63(5):723–732, 2005.
- [16] W. Reed. X-ray cargo screening systems: the technology behind image quality. *Port Technology International*, 2007.
- [17] W. Reed and E. Haines. Throughput performance factors in x-ray cargo screening systems. *Port Technology International*, 2008.
- [18] T. W. Rogers, N. Jaccard, E. J. Morton, and L. D. Griffin. Detection of cargo container loads from X-ray images. In: *Proceedings IET International Conference on Intelligent Signal Processing*, 2015.
- [19] T. W. Rogers, J. Ollier, E. J. Morton, and L. D. Griffin. Reduction of wobble artefacts in images from mobile transmission X-ray vehicle scanners. In: *Proceedings IEEE International Conference on Imaging Systems and Techniques*, 356–360, 2014.
- [20] A. Sasov, X. Liu, and P. L. Salmon. Compensation of mechanical inaccuracies in micro-CT and nano-CT. In: *Proceedings SPIE*, 70781C–70781C, 2008.
- [21] M. D. Silver, A. Sen, and S. Oishi. Determination and correction of the wobble of a C-arm gantry. In: *Proceedings SPIE*, 1459–1468, 2000.
- [22] R. Speller. Radiation-based security. *Radiation Physics and Chemistry*, 61(3-6):293–300, 2001.

- [23] H. Vogel. Search by X-rays applied technology. *European Journal of Radiology*, 63(2):227–236, 2007.
- [24] H. Vogel. Search for persons. *European Journal of Radiology*, 63(2):220–226, 2007.
- [25] H. Vogel. Vehicles, containers, railway wagons. *European Journal of Radiology*, 63(2):254–262, 2007.
- [26] H. Vogel and D. Haller. Luggage and shipped goods. *European Journal of Radiology*, 63(2):242–253, 2007.
- [27] G. Zentai. X-ray imaging for homeland security. *International Journal of Signal and Imaging Systems Engineering*, 3(1):13–20, 2010.
- [28] G. Zhang, L. Zhang, and Z. Chen. An HL curve method for material discrimination of dual energy X-ray inspection systems. In: *Proceedings IEEE Nuclear Science Symposium*, 1:326–328, 2005.
- [29] J. Zhao, X. Hu, J. Zou, G. Zhao, H. Lv, L. Xu, Y. Xu, and X. Hu. Method for correction of rotation errors in micro-CT System. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 816:149–159, 2016.